SPEEDY: An Eclipse-based IDE for invariant inference

David R. Cok
GrammaTech, Inc.
Ithaca, NY, USA
cok@frontiernet.net

Scott C. Johnson
GrammaTech, Inc.
Ithaca, NY, USA
sjohnson@grammatech.com

SPEEDY is an Eclipse-based IDE for exploring techniques that assist users in generating correct specifications, particularly including invariant inference algorithms and tools. It integrates with several back-end tools that propose invariants and will incorporate published algorithms for inferring object and loop invariants. Though the architecture is language-neutral, current SPEEDY targets C programs. Building and using SPEEDY has confirmed earlier experience demonstrating the importance of showing and editing specifications in the IDEs that developers customarily use, automating as much of the production and checking of specifications as possible, and showing counterexample information directly in the source code editing environment. As in previous work, automation of *specification checking* is provided by back-end SMT solvers. However, reducing the effort demanded of software developers using formal methods also requires a GUI design that *guides* users in writing, reviewing, and correcting specifications and automates *specification inference*.

1 Introduction

1.1 GUIs for software verification

While theoretical work on software verification technology has proceeded rapidly in recent years, there has not been a corresponding increase in using software verification for day-to-day software development. Where applied to industrial software, formal methods have remained in the hands of experts in logic and formal software verification.

A contributing cause of this problem is the lack of tools and automation that put verification tools where typical software developers need it - directly within the IDE used for other software development tasks. Furthermore, for good usability, the formal methods detail is best hidden behind a screen of automation: a work-flow that requires off-line processes or pops up windows containing interactive provers and logically stated verification tasks will not, in the authors' opinion, gain wide acceptance by developers already leery of static analysis.

We posit that the best IDEs incorporating formal methods technology must automate as many tasks as possible, including both specification discovery and specification checking; it must also present specifications and counterexample information directly in the user's working environment; where automation is not possible, specification templates and helpful advice to the user should be available. Stated differently, the GUI for software verification and the GUI for software development should be one and the same, with an integrated set of actions and features. Our vision is to create unified IDEs for managing software verification tasks, thoroughly integrated with commonly used software development tools, hiding and automating as much of the formal methods detail as possible. We see this task as combining design of user assistance features and integrating and *reducing to practice* the variety of advances in verification technology over the past decade.

This paper describes SPEEDY, a tool for specification editing, discovery, checking and debugging. Some of these elements have been demonstrated in previous work (e.g., [15], [14] and subsequent un-

Dubois C., Giannakopoulou D., Mry D. (Eds): Formal Integrated
Development Environment 2014 (F-IDE 2014)
EPTCS 149, 2014, pp. 44–57, doi:10.4204/EPTCS.149.5

© 2014. GrammaTech, Inc.

published developments), but are further integrated here. SPEEDY's particular contribution is to integrate techniques for (a) guided writing and review of specifications and (b) automated specification discovery (a.k.a. invariant inference, or function summarization)[1] into the user's working environment. Furthermore, while most work in this vein is for object-oriented languages (e.g., Java, C#, Dafny, Scala), SPEEDY targets C and C++. SPEEDY supports a development style in which top-level specifications are generated from requirements, and then refined into code and detailed specifications developed in tandem, with automated support for checking the code and specifications.

The project is funded by NASA as SBIR contracted research. This report describes the results of the Phase I research and the goals of the next phase. Thus, the tool is a work in progress; Phase I is designed to resolve technical risk and demonstrate a proof of concept. The full implementation of the prototype is a task for Phase II.

1.2 Humans and software verification

One strand of work in static analysis seeks to provide fully automated checking of software. This endeavor has seen significant advances in tools for model checking and abstract interpretation. However, the need for scalable, automated inter-procedural analysis has required approximations that limit the precision and sometimes the soundness of the resulting tools. Nevertheless, such tools are successfully distributed commercially and applied to realistic software.

However, a fully automated solution, in which the user never writes or sees a specification, is limited: it can only check the *implicit specifications* of the target software. By implicit specifications we mean those restrictions imposed by the programming language semantics, such has not dividing by zero, not accessing an array out of bounds, and not dereferencing an invalid pointer. In contrast, explicit specifications are those that state what the program is supposed to do; only with a complete, or at least substantial, set of (valid, accurate) explicit specifications can we be sure that a target program does what is intended. A useful software verification system will check that the target program is consistent with both the implicit specifications and any explicit specifications.

Where are the explicit specifications to come from? Traditionally such specifications are provided by domain experts and software verification experts, working together. This is typically a costly, labor-intensive endeavor. Thus the interest in another strand of work in program analysis: *invariant inference*. This research attempts to infer invariants about a program from the program source code. In typical specification and verification systems, the user must write specifications for each module and invariants for each loop within modules. Having automated inference techniques at one's disposal will reduce, significantly, one hopes, the effort needed to verify a set of software. One goal of the SPEEDY project is to integrate a variety of such techniques into an IDE for software development and verification; success at this integration will be a significant advance in usability of verification systems.

However, humans must stay in the software verification loop for two reasons. First, invariant inference algorithms, operating on source code, can only infer what existing software actually does, not what it is supposed to do. Thus humans are still needed to review the products of inference algorithms to check that automatically generated invariants indeed state the intent of the software. Nevertheless, it is less work to validate, correct or generalize purported specifications than it is to write specifications from scratch. Even if only the simpler 50% of specifications that are automatically generated, a substantial amount of thinking and writing by the user can be avoided. Importantly, because humans are still a part of the process, it is essential that the composite verification system have a design that promotes usability

[1] We use the terms invariant inference, specification discovery, and function summarization interchangeably.

for all levels of expertise.

Second, a software verification system can only verify that specifications and code are consistent, and not that both are correct. Careful human review processes are needed to be sure that informal requirements are correct, complete, and accurately translated into the target specification language.

Thus, SPEEDY combines GUI features for effective user review and manipulation of specifications with techniques for automatically suggesting specifications.

1.3 Contributions

SPEEDY is built on the Eclipse CDT (Eclipse's C/C++ development environment), uses back-end SMT solvers, and integrates various invariant inference tools. It is a tool for specification discovery and review, integrated into a full software and specification development environment. As in previous tools, SPEEDY provides back-end checking of specifications using SMT solvers. The key contributions of SPEEDY are these:

- it provides assistance to users in generating and reviewing specifications, building on programming productivity research by Schiller and Ernst [32];
- as one element of user assistance, SPEEDY integrates various invariant inference algorithms;
- SPEEDY provides an IDE for debugging specifications integrated with the software development environment, hiding most aspects of provers and formal methods;
- SPEEDY is built on the Eclipse CDT, providing GUI features that integrate specification manipulation with code development;
- SPEEDY is constructed with an architecture that enables evaluation of alternatives in tools and design decisions.

The following sections describe the architecture and main elements of SPEEDY and our progress in implementing our overall vision.

2 Architecture of SPEEDY

The principal goal of the SPEEDY project is to integrate, study and evaluate invariant inference tools and algorithms. However, it is also desired to assess their use in the context of actual software development. Thus SPEEDY is built as a full specification authoring, editing, discovery, checking, and debugging tool, fully integrated into an IDE for software development.

We chose Eclipse as the IDE and the Eclipse CDT (C Development Toolkit) as the C/C++ software development environment for two reasons: it was a UI environment with which we were familiar, having generated other Eclipse plugins in the past, and the contract customer desired an interface that was familiar and freely distributable. However, the more significant actions all can be invoked as command-line actions and could as well be interfaced to an alternate IDE, such as emacs or Visual Studio.

The architecture of the system is shown in Fig. 1. At the center of the architecture is a database holding internal representations of specifications in association with the corresponding source code. There are processes for reading and writing these from persistent storage; the specifications, if separate from the source code, can themselves be saved in the same version control system. Similarly the GUI interacts with the specification database to display, edit, and otherwise manipulate specifications, using a conventional model-view-controller design pattern.

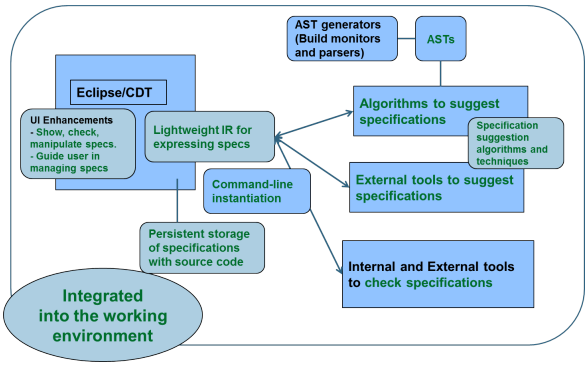
AST generators (Build monitors and parsers)
ASTs
Eclipse/CDT
UI Enhancements - Show, check, manipulate specs. - Guide user in managing specs
Lightweight IR for expressing specs
Algorithms to suggest specifications
Specification suggestion algorithms and techniques
External tools to suggest specifications
Command-line instantiation
Persistent storage of specifications with source code
Internal and External tools to check specifications
Integrated into the working environment

Figure 1: The architecture of the SPEEDY system.

The GUI itself provides mechanisms to guide users through specification writing and reviewing tasks, using Eclipse UI elements. The most significant of these is to integrate invariant inference tools and algorithms. Those also interface with the specification database through a common API, making it easy to add new inference tools and algorithms. The inference algorithms make use of sources of static analysis information in the form of abstract syntax trees attributed with type and symbol information.

Finally SPEEDY uses SMT solvers to check specifications, receiving counterexample information from the solver; the counterexample information is then displayed to the user in the editor window for the corresponding source code.

An important element of the architecture is that it be a platform for exploring various design decisions or adding new or replacement modules in the future. This is useful during the prototyping phase, but also as a research platform. Modular elements include these:

- the storage format, which is currently XML but could be any other format (section 3.2)
- the specification display format - currently we use an ACSL [6] format and experimented with VCC [13]; changing this is as easy as supplying a new parser and pretty printer (section 3.1)
- the inference tools and techniques - we integrated five different approaches, with plans for more (section 5)
- sources of ASTs - four are described in section 6
- specification checkers - we experimented with both Frama-C [24] and direct translation to SMT (section 4)
- SMT solvers - we used CVC4 [18], Z3 [19], and (through Frama-C) alt-ergo [1] - integration with Why [22] is planned for the future (section 4)

3 GUI for manipulating specifications

3.1 Specification language

At the beginning of the project we expected to spend some time assessing specification languages, and needing to choose one on which to base SPEEDY. SPEEDY needs a specification language in which it is possible to express, at least in principle, the full functionality of procedures, and not simply a relatively simple keyword language such as SAL or a C-equivalent of Java annotations. Just two languages were viable: ACSL (ANSI-C Specification Language) [6] and the VCC input language [13]. The conclusion of this investigation was that these two are similar enough that a common internal representation can be

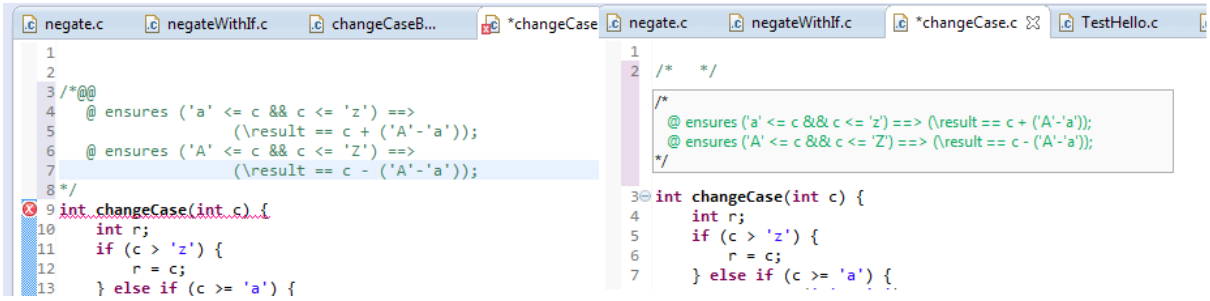

Figure 2: Two ways to weave specifications with source

used, with translation to either language as needed for display or interaction with external tools.

3.2 Storing and displaying specifications

A key design decision in a system managing code and specifications is where to store the specifications. For some languages, such as Java, the specifications can be naturally written directly in the .java files. For C, interface specifications should be associated with .h files and specifications within implementation code (such as loop invariants) will be in .c files. However, there are situations in which the source files should not or cannot be modified. The source files may not be available, such as when a library interface is being specified. Or, the development team may wish to keep the specifications separate from the source.

Accordingly, there are common use cases in which the specifications are physically separate from the source code with which they are associated. Nevertheless, to facilitate easy and accurate editing of code and specifications, the UI should present the user with the specifications directly in conjunction with the code. We call this process *weaving*: code and specifications are combined together in the editor from separate sources, and then split apart again when saved. The goal is to provide a seamless display and editing experience.

SPEEDY prototyped two ways of weaving source and specifications in the display, as shown in Figure 2. The first inserts the specifications as comments directly in the edit buffer, here matching the style of ACSL (though other styles could readily be used). The advantage is that no special graphics are needed; the disadvantage is that the line numbers change and adjustments have to be made to line numbers that reference the original source file. The second style of weaving, on the right, shows the specifications in a special, editable, styled text box inserted into the edit buffer. The text box is set to take up vertical space corresponding to its contents; thus this display style has the advantage of not changing the line numbering; the disadvantage is that careful attention to refreshing the display is required to keep the textboxes refreshed and displayed in their correct locations.

3.3 Guiding the user in creating specifications

A key element of SPEEDY is implementing mechanisms to guide the user in creating specifications. In this we are inspired by the work of Schiller and Ernst [32]. In their VeriWeb system, they measured the effects of various kinds of guidance and advice to users on the users' ability to generate and check specifications. Table 1 shows the advice mechanisms proposed by Schiller and Ernst and the corresponding implementation, with modifications and extensions, in SPEEDY.

Schiller and Ernst research result	**SPEEDY implementation**
Drag and drop editing interface	Standard Eclipse editing features: templates, code completion, context assist, keyboard shortcuts
Concrete counterexamples (from execution traces)	Concrete counterexample values and paths, from static analysis, overlaid on the source code implementation
Specification inlining, in the web tool	Two modes of specification inlining, in the source code editor
Context clues	Suggested specification locations
Specification suggestions from Daikon	Specification suggestions from multiple tools, including Daikon
Active guidance	Active guidance using integrated Eclipse Wizards, template-guided editing, and quick fixes to give suggestions

Table 1: GUI elements implemented in SPEEDY

These UI elements have two key goals. First, some of them are meant to speed up editing by incorporating mouse or keyboard actions that accomplish routine tasks. More importantly, the user's macro-level tasks and understanding are facilitated by offering automated advice and suggestions for new specifications or to correct specification errors.

3.4 Other GUI elements

SPEEDY made full use of the functionality available in Eclipse, implementing many other features that are (fairly straightforwardly) available through Eclipse extensions:

- tracking the provenance of specifications (user written or tool generated)
- a custom View to show status of specifications: Auto-generated but not yet reviewed, Accepted, Rejected, Valid, Invalid
- syntax highlighting to show status
- integration with the Eclipse build system (to launch static analysis tools)
- a custom view to show and edit specifications (as well as showing them in the editor)
- Eclipse problem markers to identify syntactic problems or failed specification checks
- showing counterexample information from failed proof attempts directly in the C editor windows, in terms of the source code (cf. section 4)
- Eclipse markers to identify locations where specifications should be inserted
- Quick fixes to insert inferred specifications in appropriate code locations
- code completion and content assists, to streamline editing
- menu items to launch various actions, including the ability to map key combinations to actions
- Eclipse preference and property sheets
- Eclipse perspective for manipulating and generating specifications
- packaging as a standard Eclipse plugin

A few more complex items were left for Phase II of the prototype development:

- Integrating the parsing (and syntax error checking) of specifications with that of code
- Syntax highlighting of keywords in specifications

- Implementing Wizards as a way to guide users through the steps of writing specifications
- Integrating specifications into the CDT's refactoring and searching tools.

Two of the above items are worth elaboration. First, SPEEDY tracks the status of specifications. As we discussed in section 1.2, a specification generated automatically may need human review – to accept it as is, to correct, to generalize, or to discard. By indicating a status – user-written, user-accepted, user-rejected, or not yet reviewed – progress of work can be recorded and remembered, even on a large scale or in a team environment. Also, changes in generated specifications or out of date specifications or checks can be highlighted.

Second, we expect to use Eclipse Wizards as a way of walking the user through the process of writing or reviewing specifications. For example, a method needs pre- and post- and frame-conditions; it is helpful to be reminded of the various pieces. Similarly loop conditions have a number of interacting aspects that are easy to omit, if not reminded of each one.

4 GUI for checking and debugging specifications

Specifications will never[2] be correct without automated means to check them; furthermore, just as with code, effective debuggers are needed to evolve specifications and code into mutual consistency.

SPEEDY uses a now common paradigm for modular checking of code and specifications (for example, see [15, 26, 23, 2]). Checking is performed procedure by procedure:

- the procedure specifications are translated into assumptions and assertions interleaved with the code, based on the semantics of the specification language. For example, preconditions become assumptions at the beginning of the procedure and postconditions are assertions at the end.
- the code, assumptions, and assertions are translated into a basic block form that uses single-assignment labeling of variables
- the basic blocks are translated into compact verification conditions (VCs)
- the verification conditions are expressed in SMTLIBv2 [4, 5] format
- an SMT solver of choice (we used primarily CVC4 [18], and demonstrated interoperability with Z3 [19]) is applied to the VC
- if the VC is invalid, a counterexample is obtained from the SMT tool
- the logical variables of the counterexample are translated back to source code variables and text locations; the values of logical variables are expressed in programming language terms
- the counterexample values and the static "execution" path are displayed in the source code editor by hover information and highlighting

Using SMTLIB allows us to plug-and-play any SMT solver, with some limitations. SMT solvers still have individual idiosyncrasies that require some tailoring of the VC structure. In addition, there are some valuable features that are missing from standard SMTLIBv2, such as conversion between Int values and Bit-Vector values, defining constant arrays, and reasoning about partially ordered sets (such as inheritance hierarchies). SPEEDY requires SMT solvers that produce counterexamples. SPEEDY also requires a combination of a number of SMT logics: Arrays, Uninterpreted Functions and either BitVectors or a combination Int and Real non-linear arithmetic. Most preferable would be a combination

[2] well, hardly ever

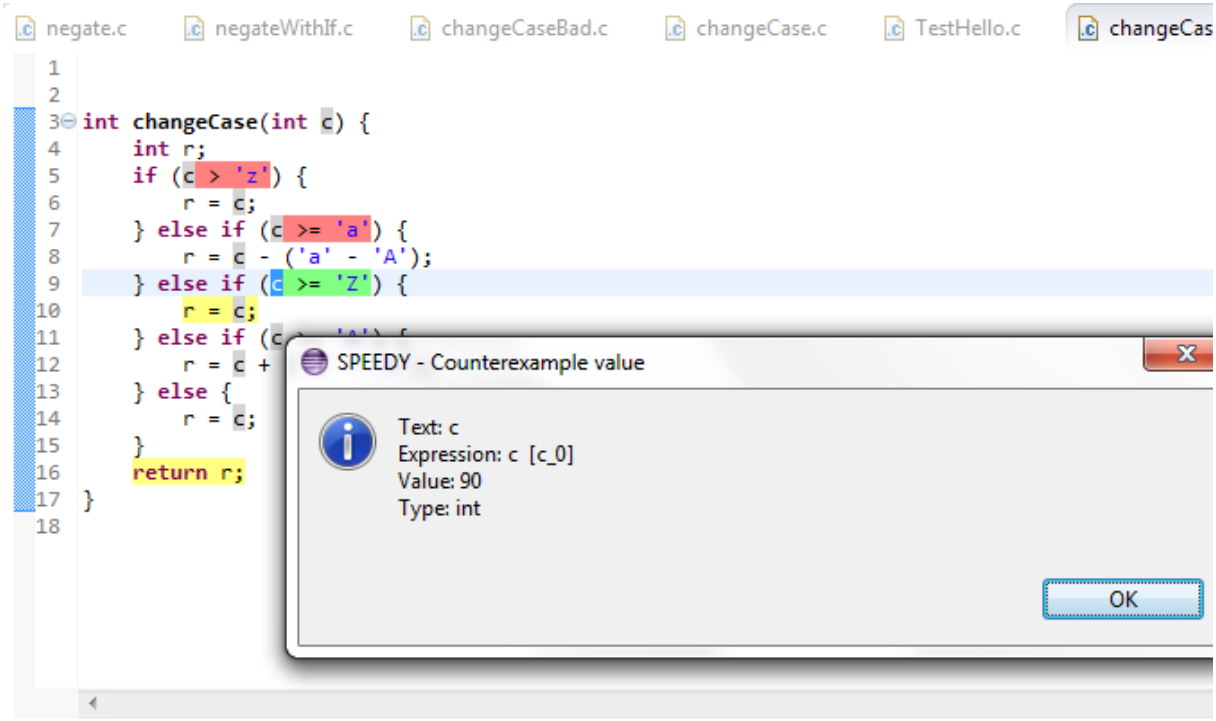

Figure 3: SPEEDY GUI showing a counterexample path and variable values

of all of the above, but not all SMT tools support this range of logics. Our demonstration can use either CVC4 or Z3.

We experimented with three techniques for verifying C code and specifications. First we combined the code and specifications into ACSL-annotated C code and checked the result using Frama-C's [24] WP plugin (which uses alt-ergo [1] as an SMT solver). This successfully verified simple procedures and specifications but does not return counterexample information. Second we performed our own conversion of C source code and its specifications to SMTLIB logical assertions, submitting the result directly to CVC4 or Z3 as SMT solvers. In this case we can obtain counterexample information from the solver, displaying it for debugging purposes as described below. A third experiment is to use Frama-C with the value analysis plugin. Of these, the second is most useful for our purposes. However, the task of creating a full translation of C and specifications to a logical encoding is substantial, a task that Frama-C already performs (for part of C, but not C++). Thus we are continuing to investigate means to obtain counterexample information through Frama-C.

From an IDE perspective, the key point is the ability to debug specifications in source code terms. Early tools, such as Simplify [20] and Esc/Java [28], provided counterexample information as an internal dump of the solver state, using many internal variable names. This information was inscrutable even to experts. The workflow above, combined with presentation in the IDE, solves this problem. This technology was adapted for C and SPEEDY from an implementation for OpenJML [15] for Java and the Java Modeling Language (JML) [10, 35].

The SPEEDY GUI for debugging specifications is shown in Fig. 3. Note that this is *static* debugging: the constraint solver seeks an assignment of values to variables that satisfies the object invariants, preconditions, background predicates, and the code itself, but falsifies one or more of the assertions,

postconditions or ending invariants of the specifications. If found, such a counterexample contains a set of inputs which, if used as the inputs in an actual execution, would trigger the identified assertions. The static debugger then allows the user to observe the control flow and variable values along the entire path, moving back and forth along the path at will.

The figure shows the execution path as highlighted statements – yellow for executed statements, red for boolean conditions (e.g., branch conditions and pre or post conditions) that are false, and green for those that are true. Variable values are shown as hover information or in dialog boxes.

5 Tools for inferring specifications

SPEEDY incorporates two ways of saving the user effort in generating specifications. One is the set of user experience design elements described above; the other is an integrated collection of tools and algorithms that directly propose invariants based on program analysis.

SPEEDY's architecture is designed to be open to existing and future specification discovery techniques. In this development phase, we integrated five mechanisms for discovering specifications. These were chosen as the first to be implemented in order to validate that the overall architecture could accommodate a variety of tools. In the next phase of development we will implement other algorithms from the verification literature. We have selected algorithms that use a variety of techniques (e.g., static analyses, dynamic analyses, symbolic execution, abstract interpretation, SMT solvers, ...) and that target various kinds of invariants (loop invariants, object invariants, preconditions, postconditions, frame conditions, memory safety, ...).

- **Specifications from CodeSonar.** CodeSonar [34] is a commercial tool for automatically finding flaws, such as buffer overruns, in a program using static analysis; it does not use specifications. Rather, it performs a bottom-up program analysis, inferring function summaries for each procedure. The function summaries include *triggers*, which are logical conditions that, if true in the calling context, indicate that a warning should be issued to the user. Thus the (negations of) triggers serve as preconditions for the procedure. SPEEDY imports the triggers from CodeSonar, expressing them in source code terms, as preconditions understandable by the user and checkable by other tools.

- **Specifications from Daikon.** Daikon [21] is a dynamic analysis tool. A front-end tool instruments the target program to record the values of various variables at various program points; then the target program is executed, perhaps over many input samples, and the trace data is collected. Daikon analyzes the trace data to find invariants that are true of all of the traces. SPEEDY then imports those invariants. Since Daikon's invariants are not necessarily true of all possible executions (just of the observed executions), these invariants will need human review. Daikon's invariants can be preconditions, postconditions, or object invariants.

- **Specifications from CodeSurfer.** CodeSurfer [12] is a static analysis tool that produces a wealth of program analysis information about a piece of software. One example is data dependencies. For this prototype, we chose to import from CodeSurfer information about the use of global variables by a procedure. This information can be presented as frame conditions.

- **Specifications from Frama-C.** Frama-C [24] is also a static analysis tool that produces program analysis data. In this case, we used Frama-C's value analysis plug-in to obtain information about the set of values a variable can have at a given program point. This information is translated into

invariants; it is also used to attest to the possibility or impossibility of various assertions being falsified.

- **Specifications by symbolic execution.** The above tools are all external tools, performing their own analyses, whose results we integrated into SPEEDY. The last integration demonstrates SPEEDY's ability to implement and integrate a specific algorithm. For the prototype, we symbolically executed the target procedure along each program path to identify preconditions, postconditions and frame conditions for the procedure. This algorithm can only handle simple procedures (loop-free, not too many branches, simple logic), but it can generate full specifications for many simple procedures, saving the user time and effort.

A point to note in this section is that it is important for a tool like SPEEDY to incorporate many kinds of specification discovery techniques: some generate preconditions, some object invariants, some loop invariants; tools will each have their own strengths and weaknesses.

6 Sources of AST, type, and static analysis information

Writing an analysis algorithm that generates specifications from source code is helped by beginning with a representation of already-parsed and type-checked program text. Thus as part of SPEEDY's design we investigated what such sources were available and usable in our context. We found four:

- Eclipse CODAN - CODAN is Eclipse's internal Code Analysis framework for C/C++. It is completely integrated in Eclipse, providing Java APIs for all the relevant parse tree (AST) and type information. We used it to implement the symbolic execution algorithm described above.
- Clang [11] - Clang's C/C++/Objective C compiler includes its own code analysis framework, intended for extension by users to provide additional static checks while compiling. Clang is an open-source, commercially-supported tool under active development, so it is readily available for use by a project like SPEEDY. Its interface is C. Clang is based on LLVM and can support analysis of any programming language for which translation to LLVM is implemented.
- CodeSurfer - GrammaTech's CodeSurfer tool also has an API for AST and program analysis information. It is also commercially released and supported. It has native C and Scheme interfaces, and through SWIG [33], many other language bindings.
- Frama-C - Frama-C [24] is an open-source static analysis environment constructed for program analysis and software verification. Many algorithms and plug-ins have been implemented on it, and it has been used for verifying safety-critical software in the European avionics industry. Its interface is ML.

One of the challenges in the next prototyping phase will be selecting among these as the foundation for algorithm implementations; we hope in fact, to be able to implement algorithms in a way that interfaces to multiple analysis infrastructures,despite the differing APIs.

7 Related work

User interfaces for several other tools have been built using Eclipse: ESC/Java2 [17], OpenJML [15, 16], and jSMTLIB [14] are examples. Similarly Spec# [3] and Dafny [27, 29] provided such functionality within Visual Studio. Leon [9], Key [7] and the Why system [22] have interfaces of their own. There are also tools primarily focused on proof management, such as PVS[30, 31], Coq [8], and Isabelle [25].

The counterexample debugging technology used here was pioneered in OpenJML [15]. OpenJML (for Java) also shares the vision of software verification technology thoroughly integrated into software development IDEs and into software developers' daily work habits. The user guidance features integrated in SPEEDY (for C) take inspiration from VeriWeb [32] (a web-based research tool for Java).

There are very many papers describing invariant inference algorithms. It is not the point of SPEEDY to develop new such algorithms, though adaptation will certainly be required. Instead, SPEEDY integrates a number of inference algorithms and techniques, with the goals of enabling easy application and comparison of these techniques. SPEEDY's innovation here is in integrating a number of such techniques into a software development IDE.

8 Conclusions and Future Work

At this point SPEEDY is a 6-month-old prototype funded by an initial contract from NASA. The Phase I work concentrated on resolving identified risks and designing an overall architecture. The prototype demonstrates all the basic UI concepts described in the paper and integrates a selection of invariant inference techniques. If the contract is renewed (or other funding is obtained), the prototype will be expanded into a deployable tool that can serve as a foundation for software verification in C, for studying invariant inference algorithms, and for evaluating designs for programming and specification productivity. Thus any conclusions based on research to date are very preliminary, but the work so far supports our key ideas:

- Eclipse, the Eclipse CDT, and the CODAN framework make a solid base for an IDE for formal methods tools
- Integrating specification manipulation features directly in the software development environment streamlines the work of creating correct specifications, even for specification experts.
- Our architecture for integrating various invariant inference techniques accommodates a wide variety of tool types.
- Eclipse's extension capabilities are adequate for the user assistance GUI features implemented or planned for SPEEDY.
- Displaying and interacting with counterexample information in the source code editor is highly effective; this design ported from Java to C without difficulty.

One intent of the project is to integrate and compare a variety of techniques; the results of that comparison are in progress and will be reported later. Another goal, however, is to create an effective IDE for working with specifications. Having a variety of techniques, with different idiosyncrasies and perhaps slightly different formal bases may be a disadvantage in creating a unified, coherent IDE and underlying specification manipulation system. The alternative is to create a fully new system or to expand one of the currently available choices. This alternative would be much more costly (in development labor). We will understand the drawbacks of the integrative approach better after the next phase of research.

The following additional work is planned:

- Study and integration of additional invariant inference algorithms
- Integration with other SMT-based constraint solvers
- Implementing additional user-guidance techniques for advising about specification generation, review, correction, and debugging
- Addition of other UI features, in particular
  - integration with Eclipse's refactoring and searching mechanisms

  - quick fix suggestions for erroneous specifications (or code)
  - extending syntax highlighting to specifications
- Scaling up the capability and set of application experiments to larger sets of code
- Comparing and evaluating invariant inference algorithms

Our work on C/C++ so far has used ACSL and the Frama-C tool. However, both of these only address a portion of ANSI-C (because they were developed to verify safety-critical avionics, which uses a 'safe' subset of ANSI-C). There is no full-fledged Behavioral Interface Specification Language for C++, similar to JML, Spec#, ACSL, or VCC. Such a language would need to combine the object-oriented features of JML and Spec# with the detailed memory buffer features contained in VCC and ACSL. Developing a C++ specification language is a significant task best addressed by the specification and verification research community as a whole.

9 Acknowledgments

The majority of this work was performed by the authors at GrammaTech, Inc., under contract to NASA (SBIR Phase I contract #NNX13CL55P). Some of the results described here are based upon work supported by the National Science Foundation under Grant No. ACI-1314674. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

10 Bibliography

References

[1] Alt-ergo: `http://alt-ergo.lri.fr/`.

[2] Mike Barnett & K. Rustan M. Leino (2005): *Weakest-precondition of unstructured programs*. In Michael D. Ernst & Thomas P. Jensen, editors: *Program Analysis For Software Tools and Engineering (PASTE)*, ACM, doi:10.1145/1108792.1108813.

[3] Mike Barnett, K. Rustan M. Leino & Wolfram Schulte (2005): *The Spec# Programming System: An Overview*. In Gilles Barthe, Lilian Burdy, Marieke Huisman, Jean-Louis Lanet & Traian Muntean, editors: *Construction and Analysis of Safe, Secure, and Interoperable Smart Devices*, *Lecture Notes in Computer Science* 3362, Springer Berlin Heidelberg, pp. 49–69, doi:10.1007/978-3-540-30569-9_3.

[4] Clark Barrett, Aaron Stump & Cesare Tinelli (2010): *The SMT-LIB Standard: Version 2.0*. Technical Report, Department of Computer Science, The University of Iowa. Available at `http://www.SMT-LIB.org`.

[5] Clark Barrett, Aaron Stump & Cesare Tinelli (2010): *The SMT-LIB Standard: Version 2.0*. In A. Gupta & D. Kroening, editors: *Proceedings of the 8th International Workshop on Satisfiability Modulo Theories (Edinburgh, England)*. Available at `http://www.smt-lib.org`.

[6] P. Baudin: *ACSL: ANSI C Specification Language*. `http://frama-c.com/download/acsl_1.4.pdf`.

[7] Bernhard Beckert, Reiner Hähnle & Peter H. Schmitt (2007): *Verification of object-oriented software : The KeY Approach*. Springer, doi:10.1007/978-3-540-69061-0.

[8] Yves Bertot & Pierre Castran (2004): *Interactive theorem proving and program development : Coq'Art : the calculus of inductive constructions*. Texts in theoretical computer science, Springer, Berlin, New York, doi:10.1007/978-3-662-07964-5. Available at `http://opac.inria.fr/record=b1101046`.

[9] Régis Blanc, Viktor Kuncak, Etienne Kneuss & Philippe Suter (2013): *An Overview of the Leon Verification System: Verification by Translation to Recursive Functions*. In: *Proceedings of the 4th Workshop on Scala*, SCALA '13, ACM, New York, NY, USA, pp. 1:1–1:10, doi:10.1145/2489837.2489838.

[10] Lilian Burdy, Yoonsik Cheon, David R. Cok, Michael D. Ernst, Joseph R. Kiniry, Gary T. Leavens, K. Rustan M. Leino & Erik Poll (2003): *An overview of JML tools and applications*. In Thomas Arts & Wan Fokkink, editors: *Eighth International Workshop on Formal Methods for Industrial Critical Systems (FMICS 03)*, *Electronic Notes in Theoretical Computer Science (ENTCS)* 80, Elsevier, pp. 73–89, doi:10.1016/S1571-0661(04)80810-7.

[11] Clang: `http://clang.llvm.org/`.

[12] CodeSurfer: `http://www.grammatech.com/research/technologies/codesurfer`.

[13] Ernie Cohen, Markus Dahlweid, Mark Hillebrand, Dirk Leinenbach, Micha Moskal, Thomas Santen, Wolfram Schulte & Stephan Tobies (2009): *VCC: A Practical System for Verifying Concurrent C*. In Stefan Berghofer, Tobias Nipkow, Christian Urban & Makarius Wenzel, editors: *Theorem Proving in Higher Order Logics*, *Lecture Notes in Computer Science* 5674, Springer Berlin Heidelberg, pp. 23–42, doi:10.1007/978-3-642-03359-9_2.

[14] David R. Cok (2011): *jSMTLIB: Tutorial, Validation and Adapter Tools for SMT-LIBv2*. In Mihaela Bobaru, Klaus Havelund, GerardJ. Holzmann & Rajeev Joshi, editors: *NASA Formal Methods*, *Lecture Notes in Computer Science* 6617, Springer Berlin Heidelberg, pp. 480–486, doi:10.1007/978-3-642-20398-5_36.

[15] David R. Cok (2011): *OpenJML: JML for Java 7 by Extending OpenJDK*. In Mihaela Bobaru, Klaus Havelund, GerardJ. Holzmann & Rajeev Joshi, editors: *NASA Formal Methods*, *Lecture Notes in Computer Science* 6617, Springer Berlin Heidelberg, pp. 472–479, doi:10.1007/978-3-642-20398-5_35.

[16] David R. Cok (2014): *OpenJML: Software verification for Java 7 using JML, OpenJDK, and Eclipse*. *Electronic Proceedings in Theoretical Computer Science (EPTCS)*. To appear.

[17] David R. Cok & Joseph R. Kiniry (2005): *ESC/Java2: Uniting ESC/Java and JML: Progress and issues in building and using ESC/Java2, including a case study involving the use of the tool to verify portions of an Internet voting tally system*. In Gilles Barthe, Lilian Burdy, Marieke Huisman, Jean-Louis Lanet & Traian Muntean, editors: *Construction and Analysis of Safe, Secure, and Interoperable Smart Devices (CASSIS 2004)*, *Lecture Notes in Computer Science* 3362, Springer-Verlag, pp. 108–128, doi:10.1007/b105030.

[18] CVC4: `http://cvc4.cs.nyu.edu`.

[19] Leonardo De Moura & Nikolaj Bjørner (2008): *Z3: An Efficient SMT Solver*. In: *Proceedings of the Theory and Practice of Software*, TACAS'08/ETAPS'08, Springer-Verlag, Berlin, Heidelberg, pp. 337–340. Available at `http://portal.acm.org/citation.cfm?id=1792734.1792766`.

[20] David Detlefs, Greg Nelson & James B. Saxe (2003): *Simplify: A Theorem Prover for Program Checking*. Technical Report HPL-2003-148, HP Labs. Available at `http://www.hpl.hp.com/techreports/2003/HPL-2003-148.pdf`.

[21] Michael D. Ernst, Jeff H. Perkins, Philip J. Guo, Stephen McCamant, Carlos Pacheco, Matthew S. Tschantz & Chen Xiao (2007): *The Daikon system for dynamic detection of likely invariants*. *Science of Computer Programming* 69(1–3), pp. 35–45, doi:10.1016/j.scico.2007.01.015.

[22] Jean-Christophe Fillitre & Andrei Paskevich (2013): *Why3 — Where Programs Meet Provers*. In Matthias Felleisen & Philippa Gardner, editors: *Programming Languages and Systems*, *Lecture Notes in Computer Science* 7792, Springer Berlin Heidelberg, pp. 125–128, doi:10.1007/978-3-642-37036-6_8.

[23] Cormac Flanagan & James B. Saxe (2001): *Avoiding exponential explosion: generating compact verification conditions*. *SIGPLAN Not.* 36(3), pp. 193–205, doi:10.1145/373243.360220.

[24] Frama-C: `http://frama-c.com/`.

[25] L. Paulson (1994): *Isabelle: A Generic Theorem Prover*. *Lecture Notes in Computer Science* 828, Springer-Verlag, doi:10.1007/BFb0030541.

[26] K. Rustan M. Leino (2005): *Efficient weakest preconditions*. *Inf. Process. Lett.* 93(6), pp. 281–288, doi:10.1016/j.ipl.2004.10.015.

[27] K. Rustan M. Leino (2010): *Dafny: An Automatic Program Verifier for Functional Correctness*. In Edmund M. Clarke & Andrei Voronkov, editors: *LPAR-16*, *Lecture Notes in Computer Science* 6355, Springer, pp. 348–370, doi:10.1007/978-3-642-17511-4_20. Available at `http://dl.acm.org/citation.cfm?id=1939141.1939161`.

[28] K. Rustan M. Leino, Greg Nelson & James B. Saxe (2000): *ESC/Java User's Manual*. Technical Note, Compaq Systems Research Center. Available at `http://www.hpl.hp.com/techreports/Compaq-DEC/SRC-TN-2000-002.pdf?jumpid=reg_R1002_USEN`.

[29] K. Rustan M. Leino & Valentin Wüstholz (2014): *The Dafny Integrated Development Environment*. *Electronic Proceedings in Theoretical Computer Science (EPTCS)*. To appear.

[30] S. Owre, J. M. Rushby, & N. Shankar (1992): *PVS: A Prototype Verification System*. In Deepak Kapur, editor: *11th International Conference on Automated Deduction (CADE)*, *Lecture Notes in Artificial Intelligence* 607, Springer-Verlag, Saratoga, NY, pp. 748–752, doi:10.1007/3-540-55602-8_217. Available at `http://www.csl.sri.com/papers/cade92-pvs/`.

[31] Sam Owre, Natarajan Shankar, John M. Rushby & David W. J. Stringer-Calvert (2001): *PVS Language Reference*. SRI International. Available at `http://pvs.csl.sri.com/doc/pvs-language-reference.pdf`. Version 2.4.

[32] Todd W. Schiller & Michael D. Ernst (2012): *Reducing the barriers to writing verified specifications*. In: *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA 2012)*, Tucson, AZ, USA, pp. 95–112. Available at `http://doi.acm.org/10.1145/2384616.2384624`.

[33] SWIG: `http://www.swig.org`.

[34] CodeSonar® web site. `http://www.grammatech.com/products/codesonar`.

[35] JML web site. `http://www.jmlspecs.org`.